

Towards Open-Vocabulary Audio-Visual Event Localization

Jinxing Zhou¹ Dan Guo² Ruohao Guo³ Yuxin Mao⁴ Jingjing Hu²
 Yiran Zhong⁵ Xiaojun Chang^{1,6} Meng Wang^{2,*}

¹MBZUAI ²HFUT ³PKU ⁴NWPU ⁵OpenNLPLab ⁶USTC

<https://github.com/jasongief/OV-AVEL>

Abstract

The Audio-Visual Event Localization (AVEL) task aims to temporally locate and classify video events that are both audible and visible. Most research in this field assumes a closed-set setting, which restricts these models' ability to handle test data containing event categories absent (unseen) during training. Recently, a few studies have explored AVEL in an open-set setting, enabling the recognition of unseen events as "unknown", but without providing category-specific semantics. In this paper, we advance the field by introducing the Open-Vocabulary Audio-Visual Event Localization (OV-AVEL) problem, which requires localizing audio-visual events and predicting explicit categories for both seen and unseen test data at inference. To address this new task, we propose the OV-AVEBench dataset, comprising 24,800 videos across 67 real-life audio-visual scenes (seen:unseen = 46:21), each with manual segment-level annotation. We also establish three evaluation metrics for this task. Moreover, we investigate two baseline approaches, one training-free and one using a further fine-tuning paradigm. Specifically, we utilize the unified multimodal space from the pretrained ImageBind model to extract audio, visual, and textual (event classes) features. The training-free baseline then determines predictions by comparing the consistency of audio-text and visual-text feature similarities. The fine-tuning baseline incorporates lightweight temporal layers to encode temporal relations within the audio and visual modalities, using OV-AVEBench training data for model fine-tuning. We evaluate these baselines on the proposed OV-AVEBench dataset and discuss potential directions for future work in this new field.

1. Introduction

Audio-visual learning, an essential sub-field of multimodal learning, has garnered increasing attention in recent years,

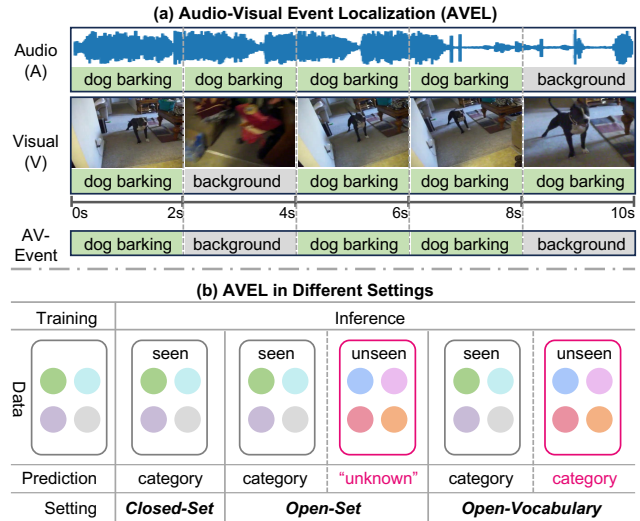


Figure 1. (a) Illustration of the AVEL task, which aims to temporally localize segments containing events that are both audible and visible, and identify their categories. (b) Studies of AVEL in different settings. In contrast to previous closed-set and open-set settings, we explore a more practical open-vocabulary AVEL problem, which needs to infer explicit event categories for both seen and unseen test data (*i.e.*, data containing classes seen and unseen during training). Each color represents a distinct event class.

resulting in the development of various research tasks, such as sound-source localization [4, 11, 13, 23, 30, 33] and segmentation [10, 12, 14, 24, 25, 50, 52], audio-visual video parsing [5, 8, 37, 45, 47, 51, 53, 54] and generation [16, 26, 35], question answering [17–20, 44], *etc.* In this paper, we focus on a fundamental research task of Audio-Visual Event Localization (AVEL) [36]. As shown in Fig. 1(a), given a video containing both audio and visual streams, the AVEL task aims to temporally localize segments that contain an audio-visual event (*i.e.*, both audible and visible) and identify its category. For segments that do not satisfy this condition (*i.e.*, only audible/visible or neither), their category is assigned to a special *background* class. In other words, this task requires perceiving the semantic alignment between audio and visual modalities.

*: Meng Wang is the corresponding author (wangmeng@hfut.edu.cn)

ties, known as audio-visual correspondence [1, 2].

In recent years, there has been rapid advancement in the AVEL field: 1) *Closed-Set Audio-Visual Event Localization*. Since the pioneering work [36], numerous significant research works have been proposed. For example, these methods aim to improve audio-visual fusion [7, 15, 21, 36, 39, 43, 48, 55], better distinguish the background [42, 49], and localize more precise temporal boundaries [22, 40, 45]. While these methods achieve satisfactory performance in traditional AVEL tasks, they are designed for a closed-set scenario. As shown in Fig. 1(b), methods in this setting can only infer data with event *classes* encountered or *seen* during model training (referred to as *seen test data* in our work), making it hard for *unseen test data* (namely test data with *classes unseen* during training) processing.

2) *Open-Set Audio-Visual Event Localization*. Recently, Yu *et al.* began exploring the AVEL task in an open-set setting [46]. To the best of our knowledge, this is currently the only work in this setting. Specifically, the open-set AVEL needs to handle both seen and unseen test data at inference. For the unseen test data with novel classes unseen in training, the model should recognize it as “*unknown*” rather than classifying it into a known category. By employing evidential deep learning and positive-unlabeled learning, [46] can identify unknown events in unseen test data. However, the model remains unable to determine specific categories for unseen events. Additionally, its model evaluation is conducted on a limited subset of the relatively small AVE [36] dataset, where only 7 classes are treated as unknown, limiting its applicability in real-world scenarios.

In this paper, we investigate the **Open-Vocabulary Audio-Visual Event Localization (OV-AVEL)** problem, a novel and more practical extension of AVEL. As shown in Fig. 1(b), OV-AVEL seeks to predict explicit event categories for both seen and unseen test data during inference, instead of assigning a general *unknown* class to unseen data as in open-set AVEL, thus providing more detailed temporal localization results. Notably, event categories in unseen test data are not present during model training. A related topic to OV-AVEL is *Audio-Visual Zero-Shot Learning (AV-ZSL)* [27–29, 31, 32], which aims to classify unseen videos during testing by integrating both audio and visual signals. The main difference is that AV-ZSL only needs to determine the category of the entire video, whereas our OV-AVEL seeks more fine-grained classification at the temporal level, requiring more precise recognition of audio-visual correspondence (*i.e.*, perceiving the event category for each modality at each segment).

To support this new task, we develop a corresponding dataset named **OV-AVEBench**. Compared to the AVE [36] dataset used in closed-set and open-set AVEL, our OV-AVEBench offers a broader range of video and event categories. An overall comparison is presented in Ta-

Table 1. **Comparison of our OV-AVEBench with other datasets used in various AVEL settings.**

Settings	Dataset	Video		Class		
		total	training	total	seen	unseen
closed-set [36]	AVE [36]	4,143	3,309	28	28	0
open-set [46]	AVE [36]	4,143	2,505	28	21	7
open-vocabulary	OV-AVEBench	24,800	13,182	67	46	21

ble 1. Specifically, the proposed OV-AVEBench includes 24,800 videos in total, approximately 6 times the number in AVE [36] dataset; The videos in our OV-AVEBench encompass 67 classes of audio-visual events, whereas the AVE dataset includes only 28. Moreover, each video sample in OV-AVEBench is manually annotated at the segment level, providing precise labels for model training or fine-tuning. These efforts in dataset construction allow us to explore more training data (seen classes) and unseen test data (unseen classes), facilitating model development and evaluation for real applications. Details about data collection, annotation, and splitting will be presented in Sec. 2.

In addition, we standardize the **evaluation metrics** for the studied OV-AVEL task. Prior AVEL studies typically adopt *accuracy* [36] as the evaluation metric, which segment-wisely compares predictions and ground truths. This metric does not account for the *recall* and may not be intuitive in evaluating predicted events across different temporal scales. Inspired by the metrics in the audio-visual video parsing task [37], we incorporate the F1-score as an additional evaluation metric for OV-AVEL, measuring it at both the segment-level and event-level. The segment-level F1-score is calculated by segment-wise comparison of predictions with ground truths. Notably, the event-level metric treats consecutive segments with identical predictions as a complete event and computes the F1-score by assessing whether the Intersection over Union (IoU) between the predicted whole event and ground truth event exceeds the threshold of 0.5. Thus, this metric is stricter in evaluating the temporal boundaries of predictions.

With the OV-AVEBench dataset and evaluation metrics established, we also propose **preliminary baselines** to address the OV-AVEL problem. To facilitate the recognition of various event classes, particularly those pertaining to unseen test data, we consider leveraging the zero-shot capability of recent language-based multimodal contrastive models. The language words are easily extendable and are not confined to predefined concepts (or categories for event classification). By applying contrastive learning to large-scale multimodal data pairs, the resulting embeddings can capture discriminative and accurate semantics. We opt to utilize ImageBind [9] because it establishes a joint embedding space across multiple modalities, aligning well with the studied OV-AVEL task. After extracting the segment-level audio, visual, and text embeddings using ImageBind (where the text represents all potential seen and un-

seen event classes), we initially explore a simple **training-free baseline**. Specifically, we compute the cosine similarity matrices for audio-text and visual-text features, respectively. In this way, we can identify the predicted event category for each audio and visual segment, corresponding to the highest similarity value, and subsequently generate audio-visual event predictions by verifying the consistency of the predicted audio and visual event categories. Notably, this baseline is training-free, directly operating on the test data. To utilize the annotated training data from the proposed OV-AVEBench dataset, we further explore a **fine-tuning baseline**. Although the unseen test data and training data possess distinct event categories, the temporal information in training data, which reflects the continuity of various audio-visual events, remains beneficial for the OV-AVEL task. Inspired by this, we incorporate some lightweight transformer layers into the ImageBind model to learn temporal relations within audio and visual modalities. Then, we fine-tune the model using the training data. Notably, parameters of the vanilla ImageBind model remain frozen, with only those of the temporal layers being learnable; thus, the increase in resource or computational load is relatively limited. Our experiments demonstrate that the fine-tuning baseline significantly outperforms the training-free version in inference on both seen and unseen test data.

In summary, our main contributions are three-fold:

- We propose the Open-Vocabulary Audio-Visual Event Localization (OV-AVEL) task, aiming to localize both seen and unseen audio-visual events in test videos. To the best of our knowledge, this work is the first to advance the AVEL area toward more practical applications in open-vocabulary scenarios.
- To facilitate this new task, we construct the OV-AVEBench dataset, which includes segment-level manual event annotations. Besides, we establish standard evaluation metrics that encompass typical accuracy, as well as segment-level and event-level F1-scores.
- We present two simple baselines: one adopting a *training-free* paradigm, which can be upgraded through further *fine-tuning* on available training data. We hope that our benchmark will inspire future research in this field.

2. OV-AVEBench

We construct the OV-AVEBench dataset to facilitate the study of Open-Vocabulary Audio-Visual Event Localization (OV-AVEL). In Table 1, we have provided a basic overview of OV-AVEBench. In the following subsections, we share more details about data collection, annotation, and splitting.

2.1. Data Collection

The target audio-visual events in the OV-AVEL task necessitate semantic correlation between the audio and visual modalities (in at least some video segments). To meet

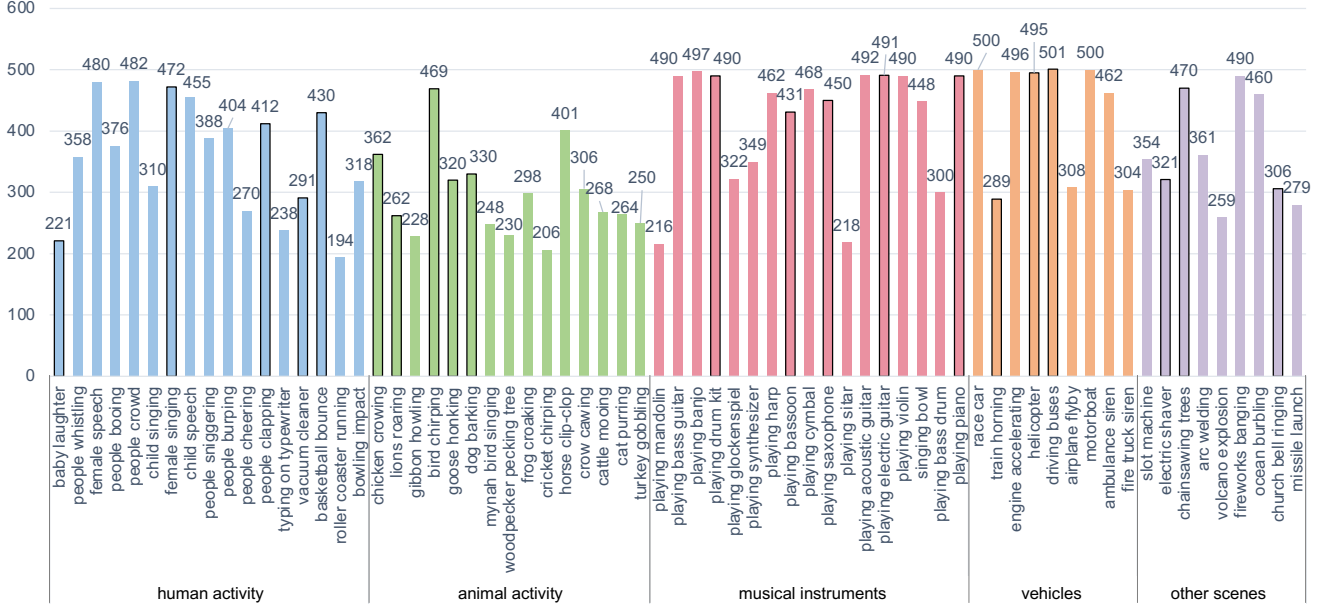
this requirement and avoid unnecessary costs, we resort to existing VGGSound [3] dataset, a large-scale audio-visual dataset in our community that provides ample video resources. Specifically, VGGSound dataset consists of over 200k videos covering 309 audio classes. However, some classes may be easily recorded in audio signal but are difficult to represent in visual frames, such as *wind noise* and *thunder*. Additionally, some classes are either too similar or too fine-grained (e.g., *car engine starting* vs. *car engine idling*), or too rare in current real-life (e.g., *dinosaurs belting*). We filtered out categories like these and ultimately selected 67 common and suitable classes for constructing the OV-AVEBench dataset. The complete category list of selected categories is presented in Fig. 2(a). These event categories correspond to five major groups: human activity, animal activity, musical instruments, vehicles, and several other audio-visual scenes from real life.

After determining the event categories, we downloaded the corresponding videos based on the YouTube URLs provided by the VGGSound dataset. A small number of videos were not currently available. Next, five volunteers were invited to manually review and check the downloaded videos. Some low-quality videos (e.g., those that were completely mismatched with their category tags) were further removed. After these steps, we ultimately retained 24,800 videos as the data resources for our OV-AVEBench dataset. The specific number of videos in each event category is shown in Fig. 2(a). Each video lasts for 10 seconds. We examine the temporal length of audio-visual events contained in the videos. As shown in Fig. 2(b), we find that 43.48% video data contain audio-visual events in all temporal segments (10 seconds), while 24.27% contain no audio-visual events (0 seconds), i.e., containing *background* class, and over one-third of the data fall in between. These video data require the model to recognize various event categories, distinguish background segments, and localize events across different temporal scales, making our OV-AVEBench dataset applicable to the OV-AVEL task.

2.2. Data Annotation

After obtaining the videos, we attempt to provide audio-visual event labels for them. For each video, we divide it into ten 1-second segments. The intermediate video frame of each segment is extracted to represent its visual component. The audio component is the corresponding 1-second audio sequence. Then, the audio-visual event labels can be determined by evaluating whether the visual frame and the audio sequence describe the same event. If they match, this segment is labeled as ‘1’ with a meaningful event category, e.g., *baby laughter*; otherwise, it is labeled as ‘0’ and categorized as *background*. This process allows us to obtain segment-level labels.

To ensure high-quality labels for the community, we con-



(a) category and video statistics of the proposed OV-AVEBench dataset




Figure 2. **Statistics about the proposed OV-AVEBench dataset.** (a) Our OV-AVEBench contains 24,800 videos covering 67 practical audio-visual scenes from the real world. Each event category and its corresponding video amount are listed. The category highlighted by a black bounding box indicates that data in that category is only available during the inference phase (unseen classes/data). (b) The audio-visual events in the videos exhibit various temporal scales, with some containing only background. We also visualize the category distribution (c), the video distribution (d) of the seen and unseen data, and the video counts for the training, validation, and test sets (e).

ducted the annotation process through crowdsourcing. Ten human annotators were involved in this process. First, we provided video examples along with guidelines to ensure that the annotators understood the annotation procedures and standards. Next, the formal annotation began. After finishing the annotation, we exchanged annotators to perform a second-round re-evaluation. Annotations with differing opinions were discussed to reach a final judgment. It took us about three weeks to complete the annotation process.

2.3. Data Splitting

The OV-AVEL task requires handling both seen and unseen data, which means that only some event categories are present during training. Our OV-AVEBench dataset contains a total of 67 event classes. As shown in Fig. 2(c), we select 46 classes as seen classes (appearing in training data) while the remaining 21 classes as unseen (only appearing

in the evaluation phase). The detailed category names can also be observed from Fig. 2(a), where the unseen classes are highlighted by black bounding boxes. Importantly, we did not simply select 21 classes from the entire category list. Instead, we identify specific classes from each major audiovisual category group while carefully balancing the number of videos in the resulting seen and unseen data. As shown in Fig. 2(d), 16,497 videos are finally used as seen data (*whose classes are seen during training*), and 8,303 videos are used as unseen data. Notably, some of the seen data can also appear in the validation and test sets. The detailed numbers are presented in Fig. 2(e). Specifically, we split the videos into training, validation, and test sets, with respective video counts of 13,182, 5,798, and 5,820. For the validation and test sets, we set the ratio of seen and unseen videos (*i.e., videos with seen and unseen classes*) at approximately 3:7. This allows us to evaluate models using more

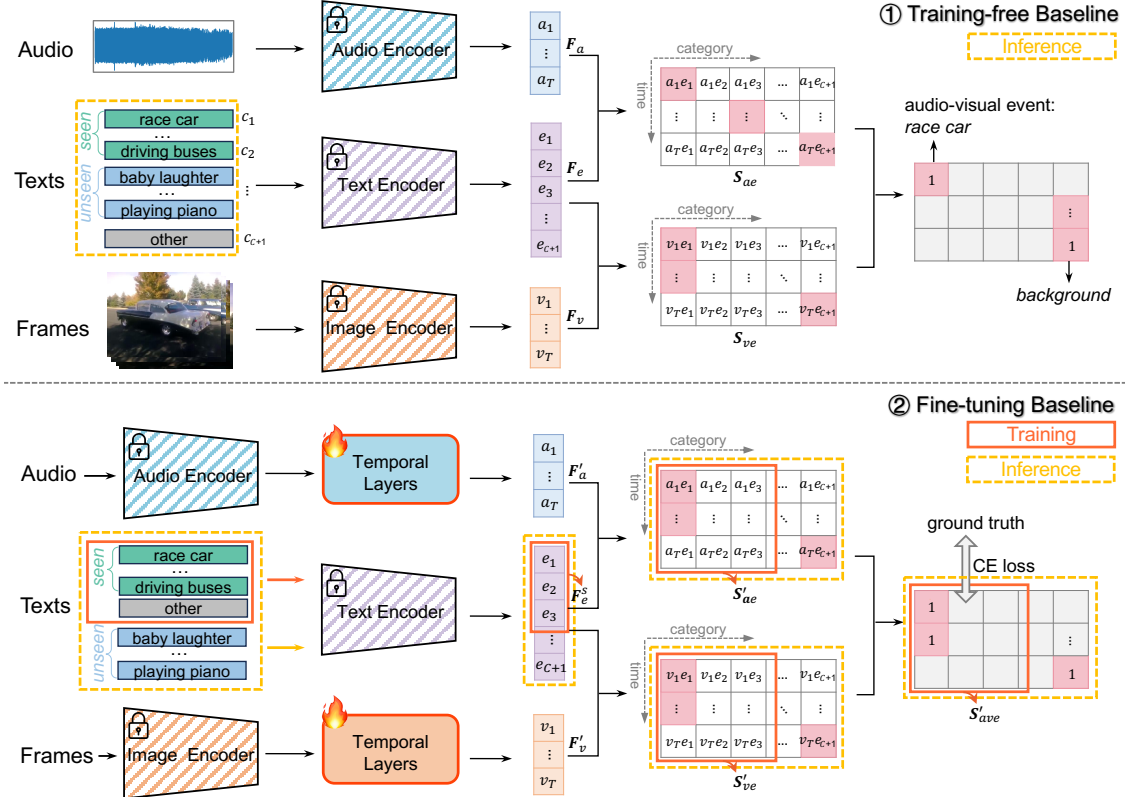


Figure 3. **Overview of the proposed baseline methods.** We utilize the audio and image encoders of the pretrained Imagebind [9] (with frozen parameters) to extract segment-level audio and visual features. ① The *training-free baseline* sends texts of all candidate classes (both seen and unseen) to extract features. Then, the audio-visual event prediction is decided by evaluating the consistency between audio-text and visual-text feature similarities. ② The *fine-tuning baseline* additionally inserts some temporal layers into the audio and visual encoders to strengthen temporal interaction learning. This model is fine-tuned/trained with training data (with seen classes). Only the texts of seen classes are known and used in model fine-tuning, while the unseen classes are additionally introduced during inference. The final audio-visual event prediction is obtained by fusing event probabilities of audio and visual modalities.

unseen data during the inference phase. We will release the OV-AVEBench dataset to the community for transparency.

3. Baselines

Task Formulation. Given an audible video, it is divided into $T = 10$ consecutive and non-overlapping segments, with $\{A_t\}_{t=1}^T$ and $\{V_t\}_{t=1}^T$ representing the audio and visual components, respectively. The OV-AVEL task aims to localize video segments that contain an audio-visual event and identify their categories. Each video is typically dominated by one event category. The ground truth labels can be denoted as $\mathbf{Y} = \{Y_t\}_{t=1}^T \in \mathbb{R}^{T \times (C+1)}$, where ‘ $C + 1$ ’ indicates the total number of audio-visual event classes plus a *background* class. Notably, during the inference phase, the OV-AVEL task addresses data with both seen and unseen classes. We denote the total number of seen and unseen event classes as C_s and C_u , respectively, where $C_s = 46$ and $C_u = 21$ in the proposed OV-AVEBench dataset. And, $C = C_s + C_u$.

3.1. A Training-free Baseline

The OV-AVEL task can be addressed with a straightforward training-free baseline, as illustrated in the upper part of Fig. 3. First, we utilize the pretrained ImageBind [9] model discussed in Sec. 1 to extract audio and visual features. Specifically, the sampled video frame from each visual segment is sent to the image encoder of ImageBind, yielding the segment-level visual features $\mathbf{F}_v = \{v_t\}_{t=1}^T \in \mathbb{R}^{T \times d}$, where $d = 1024$ is the feature dimension. Similarly, each audio segment is sent to the audio encoder to extract audio features, denoted as $\mathbf{F}_a = \{a_t\}_{t=1}^T \in \mathbb{R}^{T \times d}$.

Traditional approaches to closed-set AVEL [36] typically use only the audio and visual features as model inputs for event prediction. To achieve open-vocabulary AVEL, we adopt a zero-shot classification paradigm similar to CLIP [34]. We send all candidate event classes (seen and unseen) to the text encoder of ImageBind to obtain the text (event category) features $\mathbf{F}_e = \{e_c\}_{c=1}^{C+1} \in \mathbb{R}^{(C+1) \times d}$. Notably, we add a special text *other* that corresponds to the *background* class, handling situations that do not belong

to the listed seen and unseen classes (other new potential event classes can also be flexibly added in practical applications). Next, we compute the cosine similarities of audio-text and visual-text features, denoted as $\mathbf{S}_{ae} \in \mathbb{R}^{T \times (C+1)}$ and $\mathbf{S}_{ve} \in \mathbb{R}^{T \times (C+1)}$, as follows:

$$\mathbf{S}_{ae} = \|\mathbf{F}_a\| \otimes \|\mathbf{F}_e\|^\top, \quad \mathbf{S}_{ve} = \|\mathbf{F}_v\| \otimes \|\mathbf{F}_e\|^\top, \quad (1)$$

where $\|\cdot\|$ denotes L2-normalization and \otimes is the matrix multiplication. By scanning each row of \mathbf{S}_{ae} and \mathbf{S}_{ve} , we can predict the category of each audio and visual segment by identifying the category with the highest cosine similarity value (marked by pink boxes shown in Fig. 3①). The audio-visual events in target segments require that the category of the audio segment and the synchronized visual segment should be identical. Therefore, we can easily determine the final audio-visual event predictions by checking the audio and visual class consistency for each segment: if both modalities share the same event category, that segment contains an audio-visual event of that category; otherwise, it is recognized as *background*.

3.2. A Fine-tuning Baseline

The baseline method described above is training-free since the parameters of the audio, image, and text encoders are frozen. Since segment-level labels are available for the training data (with seen classes), we attempt to enhance the training-free baseline through additional fine-tuning.

ImageBind can provide advanced audiovisual features; however, they are independent at the segment level. The temporal relations across segments are also crucial for our OV-AVEL task because the target audio-visual events typically span various temporal scales. Motivated by this, we insert some learnable temporal layers after the audio encoder and image encoder of ImageBind to enhance the temporal interaction of each modality (illustrated in the lower part of Fig. 3). In practice, the temporal layers are implemented as a stack of L standard Transformer [38] blocks. We denote the generated audio and visual features as $\mathbf{F}'_a \in \mathbb{R}^{T \times d}$ and $\mathbf{F}'_v \in \mathbb{R}^{T \times d}$, respectively.

Training/Fine-tuning. Notably, data in the training set contains only the seen classes. Therefore, during model training/fine-tuning, only the texts of C_s seen classes and additional text *other* (outlined by the orange box in Fig. 3②) are sent to the text encoder of ImageBind to extract text features, yielding $\mathbf{F}_e^s \in \mathbb{R}^{(C_s+1) \times d}$. Then, we can compute the feature similarity matrices of audio-text and visual-text pairs, similar to Eq. 1, denoted as $\mathbf{S}'_{ae} \in \mathbb{R}^{T \times (C_s+1)}$ and $\mathbf{S}'_{ve} \in \mathbb{R}^{T \times (C_s+1)}$, respectively.

The matrices \mathbf{S}'_{ae} and \mathbf{S}'_{ve} reflect the category probabilities of audio events and visual events, respectively. We generate the final audio-visual event probability $\mathbf{S}'_{ave} \in$

$\mathbb{R}^{T \times (C_s+1)}$ by fusing them as follows:

$$\mathbf{S}'_{ave} = \sqrt{\mathbf{S}'_{ae} \odot \mathbf{S}'_{ve}}, \quad (2)$$

where \odot is the Hadamard product. This strategy differs from the direct comparison of \mathbf{S}_{ae} and \mathbf{S}_{ve} used in the training-free baseline, which is non-differentiable for model back-propagation. The ground truth $\mathbf{Y}' \in \mathbb{R}^{T \times (C_s+1)}$ for the training data can be easily obtained by selecting columns of corresponding seen classes from $\mathbf{Y} \in \mathbb{R}^{T \times (C+1)}$. Then, our fine-tuning baseline is trained by optimizing the cross entropy loss between \mathbf{S}'_{ave} and \mathbf{Y}' .

Inference. The OV-AVEL task involves handling both seen and unseen data (*i.e.*, data with seen and unseen classes) during the inference phase. As highlighted by the yellow dotted box in Fig. 3②, the texts of both seen and unseen classes are sent to the text encoder for feature extraction. The processing of audio and visual modalities follows the same flow as in training, whereas the audio and visual segments are processed by the pretrained encoders and fine-tuned temporal layers to extract audio and visual features. Then, we can generate the probability of audio-visual events by utilizing audio-text and visual-text feature similarities as described in Eq. 2. The final prediction can be made by selecting the event category with the largest probability.

4. Experiments

4.1. Implementation Details

We conduct experiments on the proposed OV-AVEBench dataset and evaluate the performance of our baselines using the three evaluation metrics introduced in Sec. 1, *i.e.*, the accuracy (Acc.), segment-level F1-score (Seg.), and event-level F1-score (Eve.). The average result of three metrics (Avg.) is also reported. In both baselines, we employ the parameters of the pretrained ImageBind_Huge¹, the only officially released version of the ImageBind [9] model, to extract audio-visual-text features. For the fine-tuning baseline, we set the batch size to 32 and fine-tune the model (learnable temporal layers) for 5 epochs; the Adam optimizer is used with a learning rate of 5e-5. All experiments are conducted on a single NVIDIA RTX 4090D (24GB) GPU. The source code will be released.

4.2. Main Results

We propose both a training-free (zero-shot) baseline and a fine-tuning baseline to address the new OV-AVEL task. Additionally, we implement several zero-shot and fine-tuning approaches to establish a comprehensive benchmark.

Comparison of Training-free approaches. We compare our training-free baseline with two methods: 1) *Video-LLaMA2* [6]. We design task-specific prompts (see our

¹<https://github.com/facebookresearch/ImageBind/tree/main>

Table 2. **Benchmark on the OV-AVEBench dataset.** We report performances of our training-free and fine-tuning baselines. Additionally, we implement and compare various training-free and fine-tuning approaches.

Method	seen				unseen				total			
	Acc.	Seg.	Eve.	Avg.	Acc.	Seg.	Eve.	Avg.	Acc.	Seg.	Eve.	Avg.
Video-LLaMA2 [6]	50.1	40.6	32.0	40.9	48.5	38.5	29.0	38.6	48.9	39.1	29.8	39.3
CLIP [34]&CLAP [41]	51.4	41.4	31.9	41.6	51.6	42.2	31.6	41.8	51.5	41.9	31.7	41.7
Training-free (our)	57.5	45.0	34.0	45.5	59.8	47.3	34.0	47.0	59.2	46.7	34.0	46.6
CMRA [43]	65.2	58.8	54.3	59.4	36.0	31.0	26.3	31.1	44.3	38.9	34.3	39.2
AVE [36]	76.6	63.6	56.0	65.4	44.6	33.2	24.0	34.0	53.8	41.9	33.2	42.9
PSP [48]	75.4	66.8	61.0	67.7	33.7	28.2	24.2	28.7	45.6	39.3	34.7	39.9
MM-Pyramid [45]	76.5	66.9	62.3	68.6	36.8	29.0	23.8	29.9	48.4	40.2	35.2	41.2
Fine-tuning (our)	72.5	61.8	54.5	62.9	64.9	55.0	47.5	55.8	67.1	56.9	49.5	57.8

supp. for details) to enable this advanced audio-visual LLM to analyze audiovisual inputs and generate event-prediction texts. 2) *CLIP&CLAP*. Instead of using ImageBind [9], we employ separate CLIP [34] and CLAP [41] models to extract audio-text and visual-text features, obtaining predictions following Eq. 1. As shown in the upper part of Table 2, our training-free baseline outperforms both methods. While Video-LLaMA2 can describe events within an entire audio sequence, it faces challenges in achieving segment-level perception of audio-visual alignment. Additionally, the comparison with CLIP&CLAP highlights the advantages of using a joint feature space for multimodal feature embedding, which better captures semantic alignment among multiple modalities for the OV-AVEL task.

Comparison of Fine-tuning approaches. We replace the temporal layers in our fine-tuning baseline with core audio-visual fusion modules from prior closed-set AVEL methods (e.g., CMRA [43], AVE [36], PSP [48], MM-Pyr [45]) to enable them perform both seen and unseen event localization for comparison. As shown in the lower part of Table 2, we find that: 1) While complex audio-visual interactions from prior methods may improve seen-class performance, they significantly degrade unseen-class performance, causing their overall results on the total test set to lag far behind ours. 2) Comparing AVE with PSP/MM-Pyr, more advanced interaction modules may exacerbate the imbalance between seen and unseen class recognition, highlighting the challenges of the OV-AVEL task.

Training-free vs. Fine-tuning baselines. Our fine-tuning baseline significantly outperforms the training-free version, showing an 11.2% improvement in the average metric (‘Avg.’) on the total test set. Moreover, we observe that: the training-free baseline model performs slightly better on the unseen test data; after fine-tuning on training data, the baseline model is improved in recognizing both seen and unseen test data (17.4%↑ and 8.8%↑ in Avg., respectively). The improvement is more pronounced for seen test data because their event classes have been seen in training. However, fine-tuning remains beneficial for unseen test data. We speculate that further fine-tuning helps the model learn tem-

Table 3. **Ablation study on the employment of the text *other*.**

Data type	w. <i>other</i>				w/o <i>other</i>			
	Acc.	Seg.	Eve.	Avg.	Acc.	Seg.	Eve.	Avg.
total	67.1	56.9	49.5	57.8	59.3	46.9	34.9	47.0
seen	72.5	61.8	54.5	62.9	62.0	49.2	37.8	49.7
unseen	64.9	55.0	47.5	55.8	58.2	46.0	33.7	45.9

poral relations from training data, facilitating the adaption and updating of prior knowledge from ImageBind to downstream OV-AVEL task. This enables more precise localization of temporal boundaries for unseen test data. This is supported by the event-level metric, which significantly improves from 34.0% to 47.5%. We provide additional evidence and discussions on this in Sec. 4.3. In short, the comparison between the two baselines highlights the benefits of further fine-tuning, especially when some training data with annotations are available. More quantitative and qualitative comparison results are provided in our *supp. material*.

4.3. Ablation Studies

In this section, we provide some ablation studies on the key configurations adopted in our fine-tuning baseline. More ablations are presented in our supplementary material.

The special class text *other*. We utilize a special text *other* to assist the model in classifying events that do not belong to either the seen or unseen classes. We conduct an ablation study to explore its impact. As shown in Table 3, the model using additional *other* class outperforms that baseline trained without *other* by 10.8% in average performance. The improvement is consistent across both seen and unseen test data. This underscores the superiority of introducing the additional class text *other*, which helps prevent the model from misclassifying unknown events or backgrounds as existing seen or unseen classes. In our *supp. material*, we further show that the employment of *other* is slightly better than other choices like *background*.

The strategy for generating S'_{ave} . In our fine-tuning baseline, we generate the audio-visual event probability S'_{ave} by computing the square root of the product of predicted audio event probability S'_{ae} and visual event probability S'_{ve}

Table 4. **Ablation study on the strategies for predicting S'_{ave} .** Detailed implementation of each strategy is shown in the main text. Results are reported on the total test data.

Strategy	Acc.	Seg.	Eve.	Avg.
Prob-avg	45.1	38.7	33.3	39.0
Fea-avg	46.8	39.8	34.0	40.2
Sqrt (Eq. 2)	67.1	56.9	49.5	57.8

Table 5. **Comparison of using temporal layers and linear layers in fine-tuning baseline.**

Data type	Temporal layer				Linear layer			
	Acc.	Seg.	Eve.	Avg.	Acc.	Seg.	Eve.	Avg.
total	67.1	56.9	49.5	57.8	46.6	38.0	32.4	39.0
seen	72.5	61.8	54.5	62.9	76.2	64.5	56.9	65.8
unseen	64.9	55.0	47.5	55.8	34.8	27.4	22.6	28.3

(Eq. 2). We refer to this strategy as **Sqrt**. Furthermore, we explore two additional variants to obtain S'_{ave} . (1) **Prob-avg** uses the average result of S'_{ae} and S'_{ve} to generate S'_{ave} , i.e., $S'_{ave} = (S'_{ae} + S'_{ve})/2$. (2) **Fea-avg** first generates the fused feature by averaging the audio feature F'_a and visual feature F'_v , and then computes the cosine similarity (sim) between the fused feature and the text feature F'_e , formulated as $S'_{ave} = \text{sim}(\frac{F'_a + F'_v}{2}, F'_e)$. We re-train the fine-tuning baseline model using these strategies and evaluate the model on the test set. As shown in Table 4, the **Sqrt** strategy significantly outperforms the **Prob-avg** and **Fea-avg** variants. The *geometric mean* used by **Sqrt** is more effective than *arithmetic mean* at preventing audiovisual event predictions from being misled by a high-probability prediction in one modality. These results indicate the importance of design in generating final audio-visual event probabilities, a factor that future research should also consider.

Temporal layer vs. Linear layer. Our fine-tuning baseline employs some *temporal layers* utilizing the self-attention mechanism in Transformer to enhance temporal interactions across video segments. Here, we replace these with learnable *linear layers* to update audio/visual features segment-wisely (i.e., without temporal interactions). As shown in Table 5, the average performance of the model fine-tuned using linear layers lags considerably behind that using temporal layers. Specifically, we find that the linear layers are slightly more effective than temporal layers for event localization of seen test data but are significantly inferior for unseen test data (27.5% \downarrow in Avg. metric). These results suggest that 1) for seen test data with classes present during training, simple linear layers may be adequate for fine-tuning; while 2) for unseen test data, sophisticated temporal relation modeling on training data becomes essential. Consequently, developing more versatile and robust network architectures would be an intriguing area for future research.

Different ratios of training data used for model fine-tuning. As shown in Table 6, we fine-tune the baseline

Table 6. **Impact of using different ratios of training data in fine-tuning baseline.**

Training	Testing	Metrics				Best epoch
		Acc.	Seg.	Eve.	Avg.	
100%	total	67.1	56.9	49.5	57.8	1
	seen	72.5	61.8	54.5	62.9	
	unseen	64.9	55.0	47.5	55.8	
75%	total	66.1	56.9	49.9	57.6	3
	seen	75.1	65.4	59.3	66.6	
	unseen	62.5	53.5	46.1	54.0	
50%	total	66.7	57.1	49.7	57.8	5
	seen	75.3	66.0	59.9	67.1	
	unseen	63.5	53.5	45.6	54.2	
25%	total	66.4	56.8	49.8	57.7	6
	seen	73.0	62.6	55.7	63.8	
	unseen	63.8	54.4	47.5	55.2	

model with various ratios of training data (sampling data for each training class accordingly). Interestingly, we find that the model achieves similar average performance across different data ratios. For instance, using only 25% of the training data, the model performance can reach 57.7% in Avg., close to that achieved with 100% training data. Additionally, we find that both 100% and 25% of training data better improve unseen test data recognition, while using 50% of training data is the most effective for seen test data recognition. These results reveal a non-linear link between training data size and the model’s performance, showing more data is not vital for seen/unseen class localization. However, the trade-off of using less training data is the need for more training or fine-tuning epochs. These findings suggest that determining a more balanced training strategy to optimize both seen and unseen data recognition would be a valuable direction for future work.

5. Conclusion

We propose the Open-Vocabulary Audio-Visual Event Localization (OV-AVEL) task, advancing the traditional closed-set AVEL problem into more practical open-vocabulary scenarios. Accordingly, we meticulously construct the OV-AVEBench dataset, making efforts in data collection, annotation, and splitting. We hope that the OV-AVEBench will serve as a standardized testbed for future research on OV-AVEL. Furthermore, we introduce two baseline approaches, a training-free baseline and a fine-tuning baseline, to address this new task. We present some discussions based on our experimental findings, which we expect will inspire future advancements in the OV-AVEL field.

Acknowledgement This work was supported in part by the National Key R&D Program of China (NO.2024YFB3311602), the National Natural Science Foundation of China (72188101, 62272144, 62020106007, and U20A20183), and the Major Project of Anhui Province (202203a05020011, 2408085J040).

References

- [1] Relja Arandjelovic and Andrew Zisserman. Look, listen and learn. In *ICCV*, pages 609–617, 2017. 2
- [2] Relja Arandjelovic and Andrew Zisserman. Objects that sound. In *ECCV*, pages 435–451, 2018. 2
- [3] Honglie Chen, Weidi Xie, Andrea Vedaldi, and Andrew Zisserman. VGGSound: A large-scale audio-visual dataset. In *ICASSP*, pages 721–725, 2020. 3
- [4] Honglie Chen, Weidi Xie, Triantafyllos Afouras, Arsha Nagrani, Andrea Vedaldi, and Andrew Zisserman. Localizing visual sounds the hard way. In *CVPR*, pages 16867–16876, 2021. 1
- [5] Yaru Chen, Ruohao Guo, Xubo Liu, Peipei Wu, Guangyao Li, Zhenbo Li, and Wenwu Wang. Cm-pie: Cross-modal perception for interactive-enhanced audio-visual video parsing. In *ICASSP*, pages 8421–8425, 2024. 1
- [6] Zesen Cheng, Sicong Leng, Hang Zhang, Yifei Xin, Xin Li, Guanzheng Chen, Yongxin Zhu, Wenqi Zhang, Ziyang Luo, Deli Zhao, et al. Videollama 2: Advancing spatial-temporal modeling and audio understanding in video-llms. *arXiv preprint arXiv:2406.07476*, 2024. 6, 7
- [7] Bin Duan, Hao Tang, Wei Wang, Ziliang Zong, Guowei Yang, and Yan Yan. Audio-visual event localization via recursive fusion by joint co-attention. In *WACV*, pages 4013–4022, 2021. 2
- [8] Junyu Gao, Mengyuan Chen, and Changsheng Xu. Collecting cross-modal presence-absence evidence for weakly-supervised audio-visual event perception. In *CVPR*, pages 18827–18836, 2023. 1
- [9] Rohit Girdhar, Alaeldin El-Nouby, Zhuang Liu, Mannat Singh, Kalyan Vasudev Alwala, Armand Joulin, and Ishan Misra. Imagebind: One embedding space to bind them all. In *CVPR*, pages 15180–15190, 2023. 2, 5, 6, 7
- [10] Ruohao Guo, Xianghua Ying, Yaru Chen, Dantong Niu, Guangyao Li, Liao Qu, Yanyu Qi, Jinxing Zhou, Bowei Xing, Wenzhen Yue, Ji Shi, Qixun Wang, Peiliang Zhang, and Buwen Liang. Audio-visual instance segmentation. *arXiv preprint arXiv:2310.18709*, 2023. 1
- [11] Ruohao Guo, Dantong Niu, Liao Qu, Yanyu Qi, Ji Shi, Wenzhen Yue, Bowei Xing, Taiyan Chen, and Xianghua Ying. Instance-level panoramic audio-visual saliency detection and ranking. In *ACM MM*, pages 9426–9434, 2024. 1
- [12] Ruohao Guo, Liao Qu, Dantong Niu, Yanyu Qi, Wenzhen Yue, Ji Shi, Bowei Xing, and Xianghua Ying. Open-vocabulary audio-visual semantic segmentation. In *ACM MM*, pages 7533–7541, 2024. 1
- [13] Ruohao Guo, Xianghua Ying, Yanyu Qi, and Liao Qu. Unitr: A unified transformer-based framework for co-object and multi-modal saliency detection. *IEEE transactions on multimedia*, 2024. 1
- [14] Dawei Hao, Yuxin Mao, Bowen He, Xiaodong Han, Yuchao Dai, and Yiran Zhong. Improving audio-visual segmentation with bidirectional generation. In *AAAI*, pages 2067–2075, 2024. 1
- [15] Xiang He, Xiangxi Liu, Yang Li, Dongcheng Zhao, Guobin Shen, Qingqun Kong, Xin Yang, and Yi Zeng. Cace-net: Co-guidance attention and contrastive enhancement for effective audio-visual event localization. *arXiv preprint arXiv:2408.01952*, 2024. 2
- [16] Bingliang Li, Fengyu Yang, Yuxin Mao, Qingwen Ye, Hongkai Chen, and Yiran Zhong. Tri-ergon: Fine-grained video-to-audio generation with multi-modal conditions and lufs control. *arXiv preprint arXiv:2412.20378*, 2024. 1
- [17] Guangyao Li, Yake Wei, Yapeng Tian, Chenliang Xu, Ji-Rong Wen, and Di Hu. Learning to answer questions in dynamic audio-visual scenarios. In *CVPR*, pages 19108–19118, 2022. 1
- [18] Guangyao Li, Wenxuan Hou, and Di Hu. Progressive spatio-temporal perception for audio-visual question answering. In *ACM MM*, pages 7808–7816, 2023.
- [19] Zhangbin Li, Dan Guo, Jinxing Zhou, Jing Zhang, and Meng Wang. Object-aware adaptive-positivity learning for audio-visual question answering. In *AAAI*, pages 3306–3314, 2024.
- [20] Zhangbin Li, Jinxing Zhou, Jing Zhang, Shengeng Tang, Kun Li, and Dan Guo. Patch-level sounding object tracking for audio-visual question answering. *arXiv preprint arXiv:2412.10749*, 2024. 1
- [21] Yan-Bo Lin, Yu-Jhe Li, and Yu-Chiang Frank Wang. Dual-modality seq2seq network for audio-visual event localization. In *ICASSP*, pages 2002–2006, 2019. 2
- [22] Tanvir Mahmud and Diana Marculescu. Ave-clip: Audioclip-based multi-window temporal transformer for audio visual event localization. In *WACV*, pages 5158–5167, 2023. 2
- [23] Tanvir Mahmud, Yapeng Tian, and Diana Marculescu. T-vsl: Text-guided visual sound source localization in mixtures. In *CVPR*, pages 26742–26751, 2024. 1
- [24] Yuxin Mao, Jing Zhang, Mochu Xiang, Yunqiu Lv, Yiran Zhong, and Yuchao Dai. Contrastive conditional latent diffusion for audio-visual segmentation. *arXiv preprint arXiv:2307.16579*, 2023. 1
- [25] Yuxin Mao, Jing Zhang, Mochu Xiang, Yiran Zhong, and Yuchao Dai. Multimodal variational auto-encoder based audio-visual segmentation. In *ICCV*, pages 954–965, 2023. 1
- [26] Yuxin Mao, Xuyang Shen, Jing Zhang, Zhen Qin, Jinxing Zhou, Mochu Xiang, Yiran Zhong, and Yuchao Dai. Tavgbench: Benchmarking text to audible-video generation. In *ACM MM*, pages 6607–6616, 2024. 1
- [27] Pratik Mazumder, Pravendra Singh, Kranti Kumar Parida, and Vinay P Nambodiri. Avgzslnet: Audio-visual generalized zero-shot learning by reconstructing label features from multi-modal embeddings. In *WACV*, pages 3090–3099, 2021. 2
- [28] Otniel-Bogdan Mercea, Thomas Hummel, A Sophia Koepke, and Zeynep Akata. Temporal and cross-modal attention for audio-visual zero-shot learning. In *ECCV*, pages 488–505. Springer, 2022.
- [29] Otniel-Bogdan Mercea, Lukas Riesch, A Koepke, and Zeynep Akata. Audio-visual generalised zero-shot learning

- with cross-modal attention and language. In *CVPR*, pages 10553–10563, 2022. 2
- [30] Shentong Mo and Pedro Morgado. Localizing visual sounds the easy way. In *ECCV*, pages 218–234. Springer, 2022. 1
- [31] Shentong Mo and Pedro Morgado. Audio-visual generalized zero-shot learning the easy way. *arXiv preprint arXiv:2407.13095*, 2024. 2
- [32] Kranti Parida, Neeraj Matiyali, Tanaya Guha, and Gaurav Sharma. Coordinated joint multimodal embeddings for generalized audio-visual zero-shot classification and retrieval of videos. In *WACV*, pages 3251–3260, 2020. 2
- [33] Rui Qian, Di Hu, Heinrich Dinkel, Mengyue Wu, Ning Xu, and Weiyao Lin. Multiple sound sources localization from coarse to fine. In *ECCV*, pages 292–308, 2020. 1
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 5, 7
- [35] Xuyang Shen, Dong Li, Jinxing Zhou, Zhen Qin, Bowen He, Xiaodong Han, Aixuan Li, Yuchao Dai, Lingpeng Kong, Meng Wang, et al. Fine-grained audible video description. In *CVPR*, pages 10585–10596, 2023. 1
- [36] Yapeng Tian, Jing Shi, Bochen Li, Zhiyao Duan, and Chenliang Xu. Audio-visual event localization in unconstrained videos. In *ECCV*, pages 247–263, 2018. 1, 2, 5, 7
- [37] Yapeng Tian, Dingzeyu Li, and Chenliang Xu. Unified multisensory perception: Weakly-supervised audio-visual video parsing. In *ECCV*, pages 436–454, 2020. 1, 2
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, pages 1–11, 2017. 6
- [39] Yu Wu, Linchao Zhu, Yan Yan, and Yi Yang. Dual attention matching for audio-visual event localization. In *ICCV*, pages 6292–6300, 2019. 2
- [40] Yiling Wu, Xinfeng Zhang, Yaowei Wang, and Qingming Huang. Span-based audio-visual localization. In *ACM MM*, pages 1252–1260, 2022. 2
- [41] Yusong Wu, Ke Chen, Tianyu Zhang, Yuchen Hui, Taylor Berg-Kirkpatrick, and Shlomo Dubnov. Large-scale contrastive language-audio pretraining with feature fusion and keyword-to-caption augmentation. In *ICASSP*, pages 1–5, 2023. 7
- [42] Yan Xia and Zhou Zhao. Cross-modal background suppression for audio-visual event localization. In *CVPR*, pages 19989–19998, 2022. 2
- [43] Haoming Xu, Runhao Zeng, Qingyao Wu, Minghui Tan, and Chuang Gan. Cross-modal relation-aware networks for audio-visual event localization. In *ACM MM*, pages 3893–3901, 2020. 2, 7
- [44] Pinci Yang, Xin Wang, Xuguang Duan, Hong Chen, Runze Hou, Cong Jin, and Wenwu Zhu. Avqa: A dataset for audio-visual question answering on videos. In *ACM MM*, pages 3480–3491, 2022. 1
- [45] Jiashuo Yu, Ying Cheng, Rui-Wei Zhao, Rui Feng, and Yuejie Zhang. MM-Pyramid: Multimodal pyramid attentional network for audio-visual event localization and video parsing. In *ACM MM*, pages 6241–6249, 2022. 1, 2, 7
- [46] Jiale Yu, Baopeng Zhang, Zhu Teng, and Jianping Fan. OpenAVE: Moving towards open set audio-visual event localization. In *ACM MM*, pages 1–10, 2024. 2
- [47] Pengcheng Zhao, Jinxing Zhou, Dan Guo, Yang Zhao, and Yanxiang Chen. Multimodal class-aware semantic enhancement network for audio-visual video parsing. *arXiv preprint arXiv:2412.11248*, 2024. 1
- [48] Jinxing Zhou, Liang Zheng, Yiran Zhong, Shijie Hao, and Meng Wang. Positive sample propagation along the audio-visual event line. In *CVPR*, pages 8436–8444, 2021. 2, 7
- [49] Jinxing Zhou, Dan Guo, and Meng Wang. Contrastive positive sample propagation along the audio-visual event line. *TPAMI*, pages 1–18, 2022. 2
- [50] Jinxing Zhou, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, and Yiran Zhong. Audio-visual segmentation. In *ECCV*, pages 386–403, 2022. 1
- [51] Jinxing Zhou, Dan Guo, Yiran Zhong, and Meng Wang. Improving audio-visual video parsing with pseudo visual labels. *arXiv preprint arXiv:2303.02344*, 2023. 1
- [52] Jinxing Zhou, Xuyang Shen, Jianyuan Wang, Jiayi Zhang, Weixuan Sun, Jing Zhang, Stan Birchfield, Dan Guo, Lingpeng Kong, Meng Wang, et al. Audio-visual segmentation with semantics. *arXiv preprint arXiv:2301.13190*, 2023. 1
- [53] Jinxing Zhou, Dan Guo, Yuxin Mao, Yiran Zhong, Xiaojun Chang, and Meng Wang. Label-anticipated event disentanglement for audio-visual video parsing. In *ECCV*, pages 1–22, 2024. 1
- [54] Jinxing Zhou, Dan Guo, Yiran Zhong, and Meng Wang. Advancing weakly-supervised audio-visual video parsing via segment-wise pseudo labeling. *IJCV*, pages 1–22, 2024. 1
- [55] Ziheng Zhou, Jinxing Zhou, Wei Qian, Shengeng Tang, Xiaojun Chang, and Dan Guo. Dense audio-visual event localization under cross-modal consistency and multi-temporal granularity collaboration. *arXiv preprint arXiv:2412.12628*, 2024. 2